

PATENT DOCKET UNI6-B130
04MV1089

Docket UNI6-B130/ 04MV1089

**METHOD AND APPARATUS FOR
AUTOMATIC EXECUTION OF CONCATENATED
METHODS ACROSS MULTIPLE
HETEROGENEOUS DATA SOURCES**

Roger Redding
Charles Hanson
Carol O'Hara
Koon-yui Poon
Thomas Johnson

Roger Redding
Charles Hanson
Carol O'Hara
Koon-yui Poon
Thomas Johnson

Title:

**METHOD AND APPARATUS FOR AUTOMATIC EXECUTION
OF CONCATENATED METHODS ACROSS
MULTIPLE HETEROGENEOUS DATA SOURCES**

Copyright Mask Work
37 CFR 1.71 (d)(e)

5

A portion of the disclosure of this patent document contains material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent disclosure, as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all copyright rights whatsoever.

FIELD OF THE INVENTION

The subject invention relates generally to data processing and more particularly to a method and apparatus providing high speed parallel accessing of data stored at a number of remote heterogeneous sites and automatic, pipelined execution 5 of successive methods on such data.

BACKGROUND OF THE INVENTION AND RELATED ART

Present technology is witnessing the development of large remote databases or "data warehouses", as well as rapid expansion of the Internet and 10 proliferation of corporate intranets. Demand is growing for increasingly large and rapid data transfers involving streaming video, visualization graphics and large data warehouse downloads over such new network protocols as the Fast Ethernet and Gigabyte Ethernet. The data which it would be desirable to access may be stored across heterogeneous sites, i.e., sites which contain different types of database 15 systems or other data containers. Hence the data which may need to be accessed may be referred to as "heterogeneous data."

At the same time as demand has grown for large and rapid data transfers, there has been constant pressure to simplify the user interface to a vast array 20 of components and data storage facilities. While individual components in a particular solution are often easy to use, combining them in a complete solution still presents extremely complex problems to the user.

In addition to simplifying the user interface to heterogeneous data and complex arrays of components, it appears desirable to provide the user with added capabilities to readily command and perform more powerful automated data 25 processing operations, in addition to simple "search" queries.

SUMMARY OF THE INVENTION

Our co-pending application, U.S. Serial No. 10/09,465,038 entitled Method
Apparatus For High Speed Parallel Accessing And Execution of Methods Across

Multiple Heterogeneous Data Sources provides the capability to access distributed data contained in a number of distributed heterogeneous data sources via a search initiated by a single JAVA Script wherein a single object represents the data to be retrieved and subjected to a method in the Script. According to the subject invention,
5 successive methods set out in a script are executed automatically across distributed data, with the results of one execution automatically piped to the next method. An example is the search of a data object followed by an automatic sort of the results of the search, followed by the e-mail of the results of the sort, wherein the data to be searched is distributed across a plurality of nodes, each node having a different type of
10 database.

The invention finds one application in a system employing metadata-based high level abstraction of a federation of clustered or distributed heterogeneous databases and/or data files in which the federation of databases is referenced or treated as a single object, as well as in an apparatus for parallel data access and concurrent
15 execution of object methods across the distributed data. The single object is referenced hereafter as the "data source object," sometimes abbreviated to simply "data object."

Still other objects, features and advantages of the present invention will become readily apparent to those skilled in the art from the following detailed
20 description, wherein is shown and described only the preferred embodiment of the invention, simply by way of illustration of the best mode contemplated of carrying out the invention. As will be realized, the invention is capable of other and different embodiments, and its several details are capable of modifications in various obvious respects, all without departing from the invention. Accordingly, the drawings and
25 description are to be regarded as illustrative in nature, and not as restrictive, and what is intended to be protected by Letters Patent is set forth in the appended claims. The present invention will become apparent when taken in conjunction with the following description and attached drawings, wherein like characters indicate like parts, and which drawings form a part of this application.

BRIEF DESCRIPTION OF THE DRAWINGS:

5 Figure 1 is a system block diagram illustrating a method and apparatus according to the preferred embodiment of the invention;

Figure 2 is a flow diagram illustrating structure and operation of an agent process according to the preferred embodiment;

10 Figure 3 is a block diagram illustrating system architecture according to the preferred embodiment;

Figure 4 is a flow diagram illustrating a messenger process according to the preferred embodiment;

Figure 5 is an inheritance diagram illustrating metadata employed according to the preferred embodiment;

15 Figure 6 is a schematic block diagram illustrating a node employing a static start-up process;

Figure 7 is a schematic block diagram illustrating a node employing a dynamic start-up process;

Figure 8 illustrates operation of an agent process at a local node in response to a request containing concatenated methods;

20 Figure 9 illustrates operation of an agent process at a remote node in response to a message generated according to Fig. 8 and;

Figures 10A-C and 11A-D illustrate respective Java Studio design panels adapted to form part of an alternate embodiment.

25 DETAILED DESCRIPTION OF ONE EMBODIMENT:

Figure 1 illustrates a plurality of remote sites or nodes 11, 13, 15, 17 wherein data to be retrieved or accessed is typically spread across the respective nodes. In the illustrative example of Figure 1, the data at node 11 comprises Microsoft NT files, the data at node 13 comprises an Oracle database, the data at node

15 comprises an SQL Server database, and the data at node 17 comprises a Microsoft Access database.

In one example of operation of the system of Figure 1, a user at user site or node 19 propounds a simple request which automatically sets in motion

5 concurrent parallel accessing of all the remote databases 11, 13, 15, 17. The request illustrated in Figure 1 is a search request and the parallel searches are referenced respectively as Search 1, Search 2, Search 3 and Search 4. The searches provide parallel access to the heterogeneous data using a metadata approach and treating the heterogeneous data as if it were a single object. The simple query or request is first
10 interpreted so as to pass the relevant part of the script from a user node across to the remote nodes. In the embodiment under discussion, queries or requests are presented as JAVA scripts.

Each of the searches is optimized with respect to the underlying data.

For example, there are number of ways of accessing the Oracle database, such as via
15 an ODBC connection or via the Oracle Call Interface. According to the preferred embodiment, the method used to access the Oracle database is via the Oracle Call Interface. This method is optimum for the purpose of the preferred embodiment because it provides the shortest path length to the data. Thus, standard database interfaces are used, while selecting the one which provides the shortest path length.
20 The user writing the query statement is unaware of the approach used to actually access the data.

The metadata describes the contents of the data object of a request (query). The metadata is contained in a repository 18, using data object models which describe the overall federation of servers and data sources. In the preferred
25 embodiment, there are four categories of data source objects:

- Distributed over the nodes of a cluster
- Distributed over a network
- Distributed over an SMP (symmetric multiprocessor)
- Not distributed

A distributed network can be an Ethernet or nodes on a cluster or a gigabit/sec connection.

A repository application generates a set of data source descriptor files automatically from the metadata at run-time. The data descriptor files contain only 5 the metadata corresponding to the data source object contained in the user-written script.

The descriptor files are held locally in NT flat files, and are used at run-time in the interpretation of the query requests. The use of optimized local files further supports high run-time performance. The repository used according to the 10 preferred embodiment is the Unisys Repository (UREP). Various other repositories could be used such as Microsoft's or a standard one such as is being developed by the Object Management Group.

The descriptor file name is also used as the name of the data object in the query scripts, which data object represents the highest level of abstraction of the 15 federation of data in question. For example, the descriptor file corresponding to an object, cluster population, would be called "cluster population." A user might write a query, for example:

cluster.population.search (if (bdate == xx/xx/xx)),
searching the population (perhaps the population of the United States) for all persons 20 with a particular birthdate. As discussed in ore detail below, an "agent" interpreting this script will refer to the local descriptor file, cluster.population, to determine the nature of the object.

In the case of Figure 1, the metadata indicates that the data is contained in the SQL Server, Oracle and/or NT files databases 11, 13, 15 and sets forth the 25 organization of all the data in the respective databases, e.g. the columns and rows and how to interpret the data stored in the database. Accordingly, the user at site 19 does not need to know the data structure and is thus writing applications at a transparent level, i.e., treating the whole network as a single object and writing methods on it.

PATENT DOCKET UNI6-BI30/04MV1089

A special interpreter or "agent" process is employed at the local or user site, which interprets the script/request and "looks up" the appropriate metadata from the NT descriptor file. The local agent then sends appropriate scripts to the particular nodes which contain data corresponding to the data object. An agent (interpreter) 5 module located at each remote node interprets and executes received scripts.

Each agent comprises a module of code (an NT process or the equivalent in another operating system). Thus, two levels of interpretation are employed, a first to interpret the script and a second to interpret and execute the interpreted script at the appropriate nodes. As much processing as possible is 10 performed close to the data, i.e., at the physical sites where the data is stored, in order to minimize message traffic between user and nodes. Thus, a function shipping model is used.

According to the example being discussed in connection with Figure 1, the agent at each remote site, 11, 13, 15, 17 receives the interpreted client request, 15 which includes a data source object name and the methods to be applied, which were originally embedded in the script generated by the user. The remote agent determines from the data source object (1) whether the data is distributed, and if so, (2) the way in which it is distributed. These details (1) and (2) are contained in the repository 18 of metadata. Once armed with items (1) and (2), the remote agent performs the required 20 method(s) upon the data.

The first level (local) interpretation of the two level interpretation process will now be further detailed in conjunction with Figure 2 and an illustrative example of operation according to the preferred embodiment of the invention. This example considers the client request as being received by an agent at the user site 19, 25 although the request could be received by an agent at a remote site.

According to step 31 of Figure 2, an agent at the user site 19 first receives the client request, which, in the preferred embodiment is in the form of a Java script. The agent at the user site 19 then interprets the script. The data source object

6042013-000060

name (e.g., C_sql_data) is embedded in the script, as are the methods to be invoked on the referenced data source (e.g., "sort" in C_sql_data.sort(state(d))).

5 The data source object is categorized by whether it is distributed, and the way in which it is distributed. The category of the data source object is specified in the data source descriptor file. As noted above, the latter is a text file with the same name as the data source object itself, i.e., C_sql_data.

10 At the beginning of the local interpretation of the script, the local agent imports the descriptor file, step 33 of Figure 2. In step 35, the local agent examines the descriptor file and determines the next processing step, depending on the category of the referenced data source object.

15 If, in step 35, the data source category is determined to be "distributed," the agent proceeds to step 37 and breaks the script into new scripts appropriate to the designated nodes. The new scripts are then sent to the designated nodes for further concurrent processing, step 38, 39 of Figure 2. The agent on the processing node checks the data source type to determine the next processing step (there are three data source types: NT file system, SQL Server, Oracle) - and then proceeds with the processing.

20 If, in step 35, the local agent determines that the data source is non-distributed, the agent proceeds to the test 41 to check to see if the data source location is local or not. If not local, the agent passes the script unchanged to the designated node, step 45, if local, the agent checks the data source type for next processing step and proceeds with processing, step 45.

25 The following code provides an example of local interpretation of the user script, C_sql_data.sort(state(d)), presented at node 1 of a cluster:

25

```
#import C_sql_data
main () {
    :
    C_sql_data.search(if (b_date == xx/xx/xx))
    :
30    Descriptor file C_sql_data resembles:
```

```
    SERVER = 1(sql_data), 3 (sql_data), 5 (nt_data)
}
Descriptor file sql_data resembles:
    SERVER = 1; MS; sql_data = publish:authors;
{
    au_id*      unique      CHARACTER(11)
    :
    State*      null        CHARACTER(2)
    :
10   }
```

According to this example, a data source object, C_sql_data, is searched for persons with a particular birthdate. A data source descriptor file, with the same name as the data source object, indicates that C_sql_data is distributed across 15 Nodes servers 1,3,5 of a cluster. Descriptor files on each node give details of the data distributed on that node (in this case, the data is in SQL Server databases on servers 1 and 3, and in an NT file system on server 5).

The agent on local server 1 begins execution of the script by importing 20 the data source descriptor file, C_sql_data. The category of the data is "cluster," the hosting server is "1" with the data distributed on servers 1,3 and 5. The agent processes the statement. In due course, the agent will check the syntax and verify, for example, that "b_date" is specified as a column in the descriptor of the sql_data object.

In processing the statement, the agent breaks the script into 25

sql_data.search()	for server 1;
sql_data. search()	for server 2;
nt_data. search()	for server 5

The agent on server 1 processes the first statement; the second statement is sent to server 3; and the third statement is sent to server 5. There is an 30 object with a descriptor file name, sql_data, on server 3 and an object with a descriptor file name nt_data on server 5. After the processing (sorting) at each node, the information is returned to the original (coordinating) agent for final processing. _

By using a function shipping model, in which the search commands are sent to be executed as close to the data as possible, and only the results ("hits") are returned to the requester, the network traffic is minimized (compared with a data shipping model, in which all the data might be sent to the requester, and the search 5 performed there). In the event that updates are involved, the approach also ensures that there will never be a later update in another server's cache, thus maintaining cache coherency across servers.

Figure 3 provides an illustrative system architecture. According to Figure 3, a Visual Basic client 51, a browser 55, or an Active Server Page, interfaces 10 to an ActiveX component 53. The client sets information to describe its request (e.g., the name of a file containing a script to be executed) in a table within the ActiveX component 53 and calls a "send" method within the component. The ActiveX component 53 interfaces with a Messenger code module 59 via a Sockets interface. In this way, the apparatus appears to the client to be an ActiveX component.

15 The "messenger" 59 listens for messages from the Sockets interface 57, and its operation is illustrated in connection with Figure 4. This module of code contains two key NT or Unix threads (or the equivalent for other operating systems): a send thread and a receive thread. The receive thread listens for new messages from a client or from an agent. The send thread returns results to the client, or sends requests 20 to another server.

As indicated by steps 63, 65, 67 of Figure 4, on receiving a message from the Sockets interface 57, the messenger 59 queues the request for interpretation by an "agent" process 61, which analyzes the message and performs the request. If, on receipt of a message, the messenger 59 detects that all agent processes are busy at 25 test 69, additional agents may be created, step 71, up to a maximum, using standard NT or Unix or equivalent operating system process initiation calls. If all agents are not busy, the next available agent process will interpret the request, as indicated by step 73.

On detecting that the data is distributed, the agent breaks the script into the appropriate scripts for each data source as discussed above and queues a request to the “messenger” process to send these scripts to the respective distributed servers to be processed in parallel. Thus, if successive “NO’s” occur at tests 65 and 75 of Figure 5 4, and a “YES” results at test 79, parallel requests are sent out. The receiving “messenger” process at the destination server queues the request to an “assistant agent” (which differs from an “agent” only in that it is always invoked from, and replies to, another “agent,” rather than to an external client). The assistant agent interprets the script (for example, a “search” of local data), queuing the results and 10 presenting a request to the local “messenger” for return to the requesting agent.

Thus, when test 83 of Figure 4 is satisfied, results are returned to the local messenger in step 84 where the results are then consolidated. The agent may then request the messenger to return results to the client, test 75, step 77. In this way, 15 automatic execution of methods is achieved across distributed heterogeneous data (in NT files, SQL server, Oracle,...) transparently to the requester without the writer of the request (script) having to be aware of where the data is located, how it is accessed, where the methods execute or how they were created. If the data is distributed, the execution runs automatically in parallel. With implementation of the agent and messenger models on different operating systems, the servers may run on a 20 heterogeneous mix of NT, Unix, 2200, A-Series, IBM,... etc.

Figure 5 is an inheritance diagram further illustrating organization of the metadata according to the embodiment under discussion. The box labeled “UREP Named Version Object” 201 represents the highest level of abstraction in the UREP and comprises a collection of data objects. The diagram of Figure 5 illustrates the 25 basic concept that each data object contains embedded data and methods (operations) applied against the data where the data further consists of attributes and types.

Figure 5 illustrates a second level of abstraction 212, which includes derived classes identified as System Node 202, System Server 203, Data Source Object 204, Field Desc 205 and System Script 206. Thus, each data object has

associated therewith information as to the system node(s) where it resides, the system servers within a node which access it, its attribute as being distributed or nondistributed, the field descriptors for NT files and the methods associated with it.

The System Node class 202 includes information sufficient to describe 5 each node in a cluster including attributes such as the Node Address which may, for example, represent an internet port sufficient to locate a node in question. The class 202 further includes construct() and destruct() methods to create or destroy a node.

The System Server class 203 includes all attributes and parameters 10 regarding each server which resides on a node, where the "server" comprises the messenger, agent and assistant agent codes, i.e., everything necessary to receive a script and to execute it. The server attribute illustrated in Figure 5 is the server port, which is the address (node and port) at which incoming messages are "listened for" by the messenger of the server in question.

The Data Source Object 204 comprises the names used for various 15 objects in the script. The attribute "DSC category" indicates whether the particular object is distributed (207) or nondistributed (208). A distributed object 207 further includes subclasses 209, 210 as to the type of distribution, i.e., across SMP nodes or across nodes of a cluster. The "ObjList" attribute gives a list of the databases contained within the distributed data source name. In other words, the object name is 20 broken down into sub-names which exist on the different nodes.

Non Distributed Data Sources 208 typically are either NT files 211 or a relational database object 213, which further break down into column, index, table and size schema 215, 216, 217, 218 as known to those skilled in the art.

The Script class 206 contains the location of any otherwise 25 unrecognized programs or methods and could contain programs or methods contained in URL's, in CORBA ORB environments, X/OPEN OLTP environments, as well as in local or remote NT executables or other script files.

Thus, a system Node contains one or more servers, each of which hosts its own set of Data Source Objects. The relationships represented in Figure 5 and

contained in the metadata indicate what Data Source Objects are related to which servers and thus supply the information necessary to create the local data source descriptor files at run-time.

The information represented by Figure 5 is preferably captured at

5 system set-up using a graphical interface under control of a system administrator with as much automation as possible in order to avoid unnecessary data entry. For example, such an interface provides automatic scanning of the rows and columns of a relational database. Once set up, the system runs applications automatically as illustrated herein.

10 The metadata may also include the location of otherwise unrecognized services, the API's (application programming interfaces) or protocols to be used in invoking services (effectively wrapping the "foreign" services). Services may also be sought in trading (OMG, ODP, etc.) networks, allowing a broad heterogeneity of service access, execution and creation. In this way, services invoked as a simple

15 JAVA method may actually have been provided in Open/OLTP, Corba objects, Microsoft DCOM/COM+, Sun EJB, Linc, MAPPER,..., or other environments. In this respect, an infrastructure is provided akin to a parallel nervous system for the invocation and integration of heterogeneous services (invoked as JAVA methods). A system according to the preferred embodiment can span platforms, OS's, and

20 architectures without a requirement for changes in the underlying OS.

In an implementation according to Figure 6, servers implementing the preferred embodiment run on all the nodes of a system which may be, for example, a cluster, a Uniysis cellular multiprocessing system (CMP), a network, or an SMP (symmetrical multiprocessor). The servers are preferably started by executing a

25 program, "jobstart," from any node in the system. "Jobstart" calls an NT service, registered as "Start Service" automatically at "boot" time on each of the systems nodes, defined in a configuration file. The "Start Service" serves as a listener on the host node in question, performing the loading and invocation of the local runtime processes comprising the messenger and agent. Multiple processes may be activated,

automatically, in the same node depending on performance considerations. As soon as the servers have been activated, the runtime process is ready to accept client requests.

In Figure 7, the configuration of Figure 6 is shown supplemented by a 5 repository (UREP). Instead of a static start-up of all the servers in the system, a dynamic invocation, based on the client (user) request, is now provided. Based on the data source name (data object) supplied in the client request, the server to which the client application is attached, in processing the user request, retrieves from the repository the details of the locations which support the data source. The Agent 10 process interpreting the scripts then dynamically activates only the servers required to support the user's request. The Agent is shown interacting with a DBMS (Database Management System). A hardware component suitable for implementing the system servers in a system like that of Figures 1, 6 or 7 is the Aquanta as manufactured by Unisys Corporation, Bluebell, Pennsylvania.

15 The Messenger is loaded and activated by the local NT service (the Start Service) on each node in the system. Initially, the client application, responding to a user's request, establishes a connection, via the WinSock interface, with this process (server). The server (process) acts as a "messenger" between the client and the agent process for the particular user. The "messenger" performs four key 20 functions:

- Acts as the "listener" to receive user requests from the client or from an agent on another node.
- Sends the results of the request back to the submitter of the request (the client or an agent on another node).
- Manages the creation of, and the assignment of tasks to, agent and assistant processes.
- Sends and receives messages to and from these agents and assistants, using shared memory.

As noted above, the Agent process accepts and sends messages from and to the request queue, maintained by the messenger. As illustrated above, the key functions performed by the agent are to parse and process each request in the JAVA script, often resulting in operations on named data sources within the system which 5 may be heterogeneous (e.g., in NT files, SQL Server, Oracle,...) and distributed. In so doing, the agent looks up the descriptor of the data source. If the data is distributed across multiple nodes, the agent rewrites the script as multiple scripts. Each of the latter scripts consists of the operations, for a particular node specified in the descriptor, to be performed on the data sets residing in that node. These scripts are 10 then sent to the "assistant" processes on these other nodes in accordance with the "function shipping" model. The system will typically be configured to run with an initial number of agent processes, with a maximum also specified.

In Figures 6 and 7 "node" is used to describe the physical hardware, e.g., an Aquanta server (as in a "node on the network" or a "node of a cluster"). A 15 server is the "apparatus" residing on that node comprising the messenger, agent and assistant code modules. Multiple servers may reside on a single node. The servers may be viewed as comprising part of a "federation." A federation is a group of servers which have access to the same data, objects and scripts. There may be more than one federation in a system.

20 Figure 8 illustrates processing of a script which contains multiple successive or "concatenated" methods. In test 31 of Figure 8, the metadata is checked by the agent at the local site to determine whether the data source is distributed. Test 31 corresponds to test 31 of Figure 2.

In Step 301, the local agent scans the script. In test 303, the local agent 25 determines whether successive methods are included in the script. If not, the routine proceeds to step 35 of Figure 2 of the original application.

If successive methods are involved, the flow proceeds to step 305 where the local agent determines which methods should be performed at the remote sites. This determination is preferably made by accessing a simple table which indicates whether

a selected method should be performed remotely adjacent the data or at the user site upon the returned results.

In step 307, the statement is broken into scripts appropriate to the servers at the remote nodes. For example, one may propound the statement:

5 population.search().sort().mail()

to search, for example, the population of the United States for people with particular attributes, sort the results of the search, and then mail the results of the sort. In such case, if the data in “population” were distributed across databases in servers on nodes 1, 3 and 5, the script:

10 population.search().sort().

is sent to the servers at each of the nodes 1, 3 and 5. Thus, in this example, the local agent has determined from a table that “search” and “sort” are methods designated for performance at the remote sites, and has generated an appropriate script to send to each of the sites.

15 The assistant agent at each of the remote servers on nodes 1, 3 and 5 then interprets the respective script and, on finding the successive methods, search().sort(), performs the first method (search()) and then leaves the results of that method stored in memory, rather than causing the results to be returned to the coordinating local agent. The second (or further) method(s) are then performed on the results of 20 the earlier method(s), and only when the results of the succession of methods are complete, are the results returned to be merged by the coordinating agent. In this way, if the data object (“population”) is distributed, the methods (search, sort) are performed automatically in parallel on the distributed data.

An example of operation of the remote agent is illustrated in Figure 9. The data object “population” 403, 405, 407 is retrieved at each of three respective nodes: 25 Node 1, Node 2 and Node 3. The method “search()” is performed by the remote agent on each respective data object, producing respective search results 409, 411, 413 stored temporarily in memory at each of the respective Nodes. The remote agent then executes “sort()” on each of the respective search results, yielding respective sort

PATENT DOCKET UNI6-BI30/04MV1089

results 415, 417, 419. The remote agents then transfer the respective sort results to the respective remote messengers, which return them to a coordinating agent at the originating site. The coordinating agent creates the merged results 421 and executes the Mail method to e-mail the final results.

5 The search, sort and mail methods are described further in connection with the following discussion of a preferred set of methods performed by the Agent and summarized in the following table:

TABLE I
SYSTEM METHODS

	<u>SQL Server</u>	<u>NT File</u>	<u>ORACLE</u>	<u>MNT</u>	<u>CLUSTER</u>	<u>TEMP</u>
1	info	info	info	info	info	info
2	format	format	format	format	format	format
3	groupby	groupby	groupby	groupby	groupby	groupby
4	compute	compute	compute	compute	compute	compute
5	search	search	search	search	search	search
6	sort	sort	sort	sort	sort	sort
7	load	load	load	load	load	load
8	copy	copy	copy	copy	copy	copy
9	extract	extract	extract	extract	extract	extract
10	remove	remove	remove	remove	remove	remove
11	modify	modify	modify	modify	modify	modify
12	join	join	join	join	join	join
13	mpl	mpl	mpl	mpl	mpl	mpl
14	size	size	size	size	size	size

The provision of such a set of key methods (the basic components of
5 every application) greatly enhances the ease of application development. Additional
methods employed are reflected in the following Table II. In addition, of course, user-
written methods may be invoked.

10

TABLE II

	<u>SQL Server</u>	<u>NT File</u>	<u>ORACLE</u>	<u>MNT</u>	<u>CLUSTER</u>	<u>TEMP</u>
1	adon	adon	adon	adon	adon	adon
2	adto	adto	adto	adto	adto	adto

PATENT DOCKET UNI6-BI30/04MV1089

3	find	find	find	find	find	find
4	update	update	update	update	update	update
5	first	first	first	first	first	first
6		mail				mail
7		print				print
8		save				save
9		saveObject				saveObject
10		SaveScript				SaveScript
11		read				read
12		write				write
13		close				close
14		index				index

Further, according to the preferred embodiment, the following logic, controls, environmental variables and commands are provided:

TABLE III
SYSTEM COMMANDS

	<u>Logic</u>	<u>control</u>	<u>environment</u>	<u>commands</u>
5	if	import	time\$	sleep ()
	else	local	date\$	close ()
	while	main	day\$	write ()
	try	private	mon\$	system ()
	catch		year\$	format ()
	continue		wday\$	audio ()
10	break		mday\$	debug ()
	dir\$	mpl ()		
15				trace ()

20 #IMPORT: control

The import control is used to identify a data object from the repository. The object name following the import statement is a unique name in the system. As noted, the different data sources processed according to the preferred embodiment are SQL server, Oracle, NT files, and Multiply NT files. The data object may, of course, include other data sources.

The import statement makes available to the script interpreter (agent) all the metadata files describing the contents of the selected data object. The first occurrence of "import" causes the appropriate data source descriptor file to be set up at each node of the system. Each such descriptor file containing the metadata

corresponding to the data source object in the user written script. The API required to access the data is also determined within this statement. The developer never has to be concerned about the type of data, whether the data is local, clustered, or even if the data is moved or columns modified. Because the script is interpreted at run time, all current column sizes, data location, API types, whether to use parallel methods and etc. are all handled dynamically.

Example:

```
10          #import    Personnel
              main ()      {
                  Personnel.sort ( birthday ) ;
              } / /main
```

In the above example, ...

15

#LOCAL: control

This control is used to identify a temporary data object or record set. The object name following the local statement will be a unique name for this script. The different temporary data sources which can be processed are SQL server, Oracle, NT files, clustered data sources and multiple NT files.

The LOCAL statement makes all connections required for this data object. It is possible for a data object to consist of Oracle and SQL server or any other data source. If there are multiple data sources, all connections are handled by this statement. The API required to access the data is also determined within this statement. The developer never has to be concerned about the type of data, whether the data is local, clustered, or even if the data is moved or columns modified. Because the script is interpreted at run time all current column sizes, data location, API types whether to use parallel methods, etc., are all handled dynamically.

Example:

```
#import          Person
#local    recset    Result = new Person;
#local    recset    Rslt2 = {Person.first_name  Person.last_name}
5      #local    nt      picture = @data/bitmaps/picture.bmp{}
#local    ms      table = data_source:tablename{1_name character(10)
f_name character (15) ssn character(9)}

10      main ()  {
        Result.load(Person);
        Rslt2.load(Result);
    } //main
```

15 **DATA OBJECT IDENTIFIERS:**

The following data object identifiers are used on the local control.
This allows the interpreter to know which API to use to reference the data object.

	IDENTIFIER	DATA API
20	MS	Microsoft SQL Server
	ORA	Oracle
	NT	NT Files
	RECSET	Temporary Table

SCRIPT VARIABLE TYPES:

The following script variable types are supported:

CHAR
INTEGER
5 SHORT
LONG
FLOAT
REAL
DOUBLE
10 DECIMAL
NUMERIC
BYTE
STRING
RECORD
15 DATE
TIME
TIMESTAMP

SQL DATA TYPES:

20 The following SQL data types in record databases are supported:

	<u>SQL type</u>	<u>Size</u>	<u>Data Type</u>
	SQL_CHAR	1	char
	SQL_VARCHAR	1	char
25	SQL_BIT	1	char
	SQL_TINYINT	4	long
	SQL_SMALLINT	4	long
	SQL_INTEGER	4	long
	SQL_BITINT	4	long
30	SQL_REAL	8	double
	SQL_FLOAT	8	double
	SQL_DECIMAL	8	double
	SQL_NUMERIC	8	double
	SQL_DOUBLE	8	double

PATENT DOCKET UNI6-BI30/04MV1089

SQL_DATE	6	DATE_STRUCT
SQL_TIME	6	TIME_STRUCT
SQL_TIMESTAMP	16	TIMESTAMP_STRUCT

5

PRIVATE: control

This control identifies and creates a variable for this script.

Here are some samples:

10	private	int	aa;
	private	long	aa;
	private	short	aa;
	private	int	aa,bb;
	private	int	aa = 10, bb = 20;
15	private	string	strg;
	private	char	chr;
	private	char[20]	buf;
	private	char	buf [201] = "a b c d e f g";
	private	double	dbl = 10.25;
20	private	float	flt = 10.25;
	private	record	rec = new data_object

ENVIRONMENTAL VARIABLES:

25 The system environmental variables can be used in the script just as any other string variable. There are also additional string variables and two reserved words, they are listed below:

time\$	//	current time
--------	----	--------------

PATENT DOCKET UNI6-BI30/04MV1089

	date\$	//	current date
	day\$	//	current day
	mon\$	//	current month
	year\$	//	current year
5	wday\$	//	current week day
	mday\$	//	current month day
	dir\$	//	base director for this federation

10 Reserved words:

TRUE	//	A non zero value
FALSE	//	A zero value

15 **#TRACE: control**
This control will activate the trace code to aid in the debug of a script. It will write script logic records to the script trace file as they are processed. The write command will also write data to the same trace file. The greater the number, the more information dumped to the trace file. It is presently preferred to either use a two or three only.

An example of the "trace" control is:

```
5      #trace    2      //dump script record to trace file
          //before it is executed
      #import          Personnel
      main()  {
          Personnel.search(
              if      (birthday == "06/14/45");
          );
10     write("Search complete");
      close().mail(Charles.Hanson@unisys.com, "trace");
      }  //main
```

NOTE: The code in the "box" above identifies embedded script code.

15 The embedded script code is contained as a parameter within the relevant method and will be interpreted as part of the definition of what the particular method should perform.

FORMAT(): function

20 This function is used to create a character string from multiple arguments. The syntax and rules of sprintf apply.

An example of the 'format' function is:

```
25      private      char      buf[20];
      private      int       cnt = 25;
      private      char [20]   name = "total";
      main()  {
          buf = format ("%s count=%d",name,cnt);
      }  //main
```

UPPER(): function

This function is used to convert the argument character string to upper case characters. It does not change the argument variable.

5

An example of the 'upper' function is:

```
10      private      char          buf[20]
          private      char[20]        name = "abcdef";
main()  {
          buf = upper( name);
} // main
```

15

LOWER(): function

This function is used to convert the argument character string to lower case characters. It does not change the argument variable.

An example of the 'lower' function is:

```
20      private      char          buf[20];
          private      string        name = "ABCDEFGHI";
main()  {
          buf = lower( name);
} // main
```

25

STRIP(): function

This function is used to strip off all leading and trailing spaces in the argument character string. It does not change the argument variable.

An example is:

```
5      private      char      buf[20];
      private      char[20]    name = "  abcdef  ";
      main()  {
          buf = strip( name);
      } //main
```

10

CENTER(): function

This function is used to center the character string in the argument. It does not change the variable and the variable must be a character type of fixed length.

An example of the 'center' function is:

```
15     private      char      buf[20];
      private      char[20]    name = "abcdef";
      main()  {
          buf = center( name);
      } //main
```

20

LEFT (): function

This function is used to left justify the character string in the argument. It does not change the variable and the variable must be a character type of fixed length.

An example is:

```
15      private      char      buf[20];
      private      char[20]    name = "      abcdef";
      main()  {
      5       buf = left(name);
      } // main
```

RIGHT(): function

10 This function is used to right justify the character string in the argument. It does not change the variable and the variable must be a character type of fixed length.

An example of the 'right' function is:

```
15      private      char      buf[20];
      private      char[20]    name = "abcdef";
      main()  {
      15     buf = right(name);
      } //main
```

20

FOR(): function

This function is the same as the 'for' function in JAVA.

An example of the 'for' function is:

```
25      private      int      a;
      main()  {
      25     for( a=0; a < 10; ++a)  {
      } // do some logic
      } ;
      } // main
```

WHILE(): function

This function is the same as the 'while' function in JAVA.

An example of the 'while' function is:

```
5      private      int      a = 0;  
      main()  {  
      while(a++ < 10)  {  
          // do some logic  
      } //while  
  10    } // main
```

BREAK(): function

This function is the same as the 'break' function in JAVA.

15 An example of the 'break' function is:

```
      private      int      a;  
      main()  {  
      for( a=0; a < 10; ++a)  {  
          if ( a == 5)  {  
              20        break  
          }  
          // do some logic  
      } // for  
  25    } // main
```

CONTINUE(): function

This function is the same as the 'continue' function in JAVA.

An example of the 'continue' function is:

```
private           int           a;  
5  
main()  {  
    for( a=0; a < 10; ++a)  {  
        if ( a == 8)  {  
            a = 0  
            continue  
10          }  
            //  do some logic  
        } // for  
    } // main
```

15

EXIT(): function

This function is the same as the 'exit' function in JAVA.

An example of the 'exit' function is:

```
main()  {  
20    if ( more_to_do == FALSE)  {  
        exit();  
    }  
} // main
```

25

IF(): function

This function is the same as the 'if' function in JAVA.

An example of the 'if' function is:

```
5      private      int      a = 0;  
      main()  {  
      if ( a < 10 ) {  
          // do some logic  
      } else {  
          // do some more logic  
      }  
10     }  
      } // main
```

TRY: function

15 This function will allow you to watch for application errors and then
break out of the code and jump to a catch routine. It is the same as the 'try' function
in JAVA.

An example of the 'try' function is:

```
20     main()  {  
      try  {  
          // do some logic  
      } catch () {  
          // do error handling  
      }  
25     } // main
```

CATCH: function

This function is for handling error conditions. The argument of the control is ignored currently.

An example of the 'catch' function is:

```
5      main()  {
          try  {
              // do some logic
          } catch () {
              // do error handling
          }
      } // main
```

SWITCH: function

15 This function allows you to selectively do code depending on the value of a variable. It works in conjunction with case statement. It is the same as the 'switch' function in JAVA.

An example of the 'switch' function is:

```
20      private int      cnt = 1
          main()  {
              switch (cnt)  {
                  case 0:
                      break;
                  case 1:
                      break;
                  default:
                      break;
              } //switch
          } // main
```

CASE: function

This function allows you to selectively do code when the value of the switch statement object matches the value of the case statement. It works in 5 conjunction with the switch statement.. It is the same as the 'case' statement in JAVA.

An example of the 'case' function is:

```
private int      cnt = 1
main()  {
    switch (cnt)  {
        case 0:
            break;
        case 1:
            break;
        15      default:
            break;
    } //switch
} // main
```

20

DEFAULT: function

This function will allow you to identify code as default when there is not a match for the case statement. It works in conjunction with switch and case statements. It is the same as the 'default' statement in JAVA.

An example of the 'default' function is:

```
private int      cnt = 1
main()  {
    switch (cnt)  {
        5      case 0:
                break;
        case 1:
                break;
        default:
                break;
    10     } //switch
} // main
```

15 SQLINFO: (SQL SERVER) command

This command creates a temporary data object that contains the information included in the data source name and table name referenced in the argument. The information contained in the data resulting file identifies the server where the data object is located, the type of data (SQL SERVER) and its data source name and table name. It also includes all column names, types and sizes. The object created by this command may then be added to the system with the SAVEOBJECT method discussed in further detail below.

The following example illustrates connection to a SQL SERVER data source with the name of "tpcd" to create the metadata for the table "lineitem." This information would then be recorded as a data object ("tpcd1") in the metadata. This is to make it easy to create the metadata corresponding to existing SQL Server and Oracle files.

Example:

```
main() {  
    sqlinfo ("tpcd:lineitem") .saveobject ("tpcd1") ;  
} //main
```

5

ORACLEINFO: (ORACLE) command

This command creates a temporary data object that contains the information included in the data source name and table name specified in the 10 argument. The information contained in the resulting file identifies the server where the data object is located, the type of data (ORACLE) and its data source name and table name. It also includes all column names, types and sizes. The object created by this command could then be added to the system with the saveobject method.

15 The example connects to an Oracle data source with the name of "tpcd" and creates the metadata for the table "lineitem". This information is then recorded as a data object ("tpcd2") in the metadata.

Example.

```
main() {  
    oracleinfo("tpcd:lineitem").saveobject("tpcd2");  
} //main
```

60 61 62 63 64 65 66 67 68 69 70

MAIL: command

This command is used to send e-mail. There are three arguments of which only the first one is required. The first argument is a character string containing the person's e-mail address. The second is the subject and the third 5 argument is the body of the message.

The example below searches the Personnel data object for employees with a birthday today identified by environmental variable date\$. It then reads the records and sends an e-mail to each employee with a birthday.

Example:

```
10      #import      Personnel
          private record  rec = new Personnel;
          private char     buf[200];
          private string   bf="hope you have a wonderful day"
          Personnel.search(
15          if (birthday == "date$);
          );
          while (rec = this.read() )  {
          buf = format ("Happy birthday %s %s",
                        rec.first_name,bf);
20          mail (rec.email, "birthday greetings",buf);
          } // while
          } // main.
```

WRITE: command

This command is used to aid in the debug of applications. The data specified is written to the trace file of the script. The trace command also writes to this same file.

5 There are different forms to the arguments.

```
write("character string only");  
write(data_record_object);
```

The example below will write the character string with the value of cnt to the trace file.

```
10  private int      cnt = 25  
      main()  {  
          write(format("The value of cnt = %d", cnt ))  
      } // main
```

15

SLEEP: command

This command will suspend the script for the amount of milliseconds specified in the argument.

The example below will suspend the script for one minute.

20 Example:

```
main()  {  
    sleep (60000);  
} // main
```

25

CLOSE: command

This command will close the trace file and make it the current "this" object.

5 This example creates a trace file and e-mails it to a fictitious programmer.

```
#trace 2
#import Personnel
main() {
    Personnel.search(
10    if (birthday == "06/14/55");
    );
    close();
    this.mail("Charles.Hanson@unisys.com," , "ERR");
} //main
15
```

OBJECTS: command:

20 This command will create a temporary data object that contains a list of all the data object names in the federation.

The example will display using notepad a list of all data objects in this federation.

```
main() {
    objects().display();
25} //main
```

FORMAT: command

This command will format a character string. The basic rules of a sprintf command in c++ apply.

The example will write the current date and time.

5

Example:

10

```
#import      Client
private char  buf[100];

main()  {
    buf = format("%s - - %s", date$,time$);
    Client.write(buf);
} //main
```

15

MPL: command

This command starts either a named script or a data object, which contains a script, on the server specified by the second argument. If the second argument is omitted then it will start the script on the local server. The new started 20 script runs as an independent script and will not return to the requesting script.

Example:

```
#local view script={ }
```

```
main ()  {  
5      mpl("strip_data",4);  
      script.write("#import Personnel      ");  
      script.write("main ()  {      ");  
      script.write("Personnel.search(      ");  
      script.write("  if(state== /"CA/");      ");  
10     script.write(")      ");  
      script.write("this.format(/%s,%s/",      ");  
      script.write("last_name, first_name) ;      ");  
      script.write("this.print() ;      ");  
      script.write("}      ");  
15     script.close();  
      mpl(script,5);  
}  //main
```

EXECUTE: command

20 This command executes the program of the first argument on the server specified by the second argument. If the second argument is omitted, then it will start the executable program on the local server. The executable program runs as an independent program and will not return to the requesting script.

25 Example:

```
main()  {  
      execute("xyz.exe",4);  
      execute("xyz.bat",5);  
}  //main
```

AUDIO: command

This command will play the wav file specified by the argument. The wav file must exist in the directory specified for messages.

5

Example:

```
main() {  
    audio("greetings");  
} //main
```

10

WRITE: method

This method is used to send multiple character records to the data object referenced in the script. The method supplies one record at a time.

15

```
Client.write("character string only");  
Client.write(data_object);
```

20

The example below searches the Personnel data object for employees with a birthday on June 14. Notice the wild card character ('**X**') is used to only search part of the column. The result of the search will be reformatted into three columns with commas separating them.

Example:

```
#import Personnel
#import Client
main() {
5   Personnel.search (
        if (birthday == "06/14/ |");
    )
    this.format("%s,%s,%s",last_name,first_name,ssn);
    Client.write(this);
10  } //
```

WRITE: method of temporary data objects

This method is used for constructing reports that one would e-mail,
15 fax, print or send to a bean.

There are different forms to the arguments.

```
Data_object.write("character string only");
Data_object.write(data_record_object);
```

The first use of this method in a script creates a new object or erases
20 the existing object if it existed. Every reference after the first adds records to the
method's object. Before other methods can reference this new object a 'close' method
must be performed on it.

The example below writes two records to the temporary data object
called temp and then sends the result to the bean. Notice that the close method is
25 required before the data is sent. The second example below achieves the same results.

Example:

```
#import          Client
#local          recset      temp={}
main() {
5           temp.write(format(The number of records = %d",
25 )); 
           temp.write (" This will be the second record of the data object called
temp")
10          temp.close;
           Client.write(temp);
} //
```

INFO: method

15 This method creates a temporary NT file that contains the data definition for the method's object. The information contained in the data definition file identifies the server where the data object is located, the type of data (Oracle, SQL SERVER,NT file) and its table name and data space. It also includes all column names, types and sizes.

20 Example:

```
#import          Personnel
main() {
25           Personnel.info().display();
} //main
```

FORMAT: method

This method creates a temporary data object by merging character strings and data columns from the data object. The basic rules of a `sprintf` command in C++ apply. Every record of the method's object is processed.

5 The example below creates a temporary data object. The record of the new object contains the data from two columns (last_name, first_name) separated with two '-' characters. The temporary data object will have the same number of records as the original data object and is sent to a bean.

Example:

```
10          #import Personnel
          #import Client

          main()  {
15          Personnel.format("%s -- %s",
                           last_name, first_name);
          Client.write(this);
          } //main.
```

20 **LOAD:** method

This method loads the method's object with the data object in its arguments. If the column names of the objects do not match, the columns are left blank. Only the like columns are loaded. If columns are the same type but different lengths they are either truncated or expanded to the size of the receiving column. If the data object already existed all records will be deleted, and the data object will only contain the new records. If the data object did not exist, a new one will be created. If the column types changed from the previous data object, then a remove method must be called before overloading the new object.

Example:

```
#import          Personnel
#import          Payroll
```

```
5      main() {
          Payroll.load(Personnel);
      } //mail.
```

10 SORT: method

This method sorts the method's object by the column(s) in its arguments. If there are multiple columns, the order of the sort is determined by the order of the columns. Descending sort is specified by a "(d)" following the column name. If the arguments contain a numeric value, that value determines the maximum 15 number of records to return.

The example below sorts the Personnel data object on state column (ascending) and the secondary level sort on last_name column in descending order. The numeric "10" specifies to return a maximum of 10 records. A maximum of 10 records containing two columns (state and last_name) will be returned.

20

Example:

```
#import      Client
# import      Personnel
#local  recset temp={ Personnel,state
5          Personnel.last_name }

main() {
    Personnel.sort( state, last_name(d),10 );
    temp.load(this);
    Client.write(temp);
10
} /main
```

SEARCH: method

This method uses an embedded script within its arguments. The
15 embedded script is interpreted by the method and executed for every record in the
method's object. The result of this method is a temporary data object containing all
the records that match the search criteria.

The example below searches the Personnel data object and sorts all
records where the state column is either "CA" or "MN". The records returned contain
20 two columns (state and last name) that are separated by a comma.

Example:

```
#import          Client
#import          Personnel

5      main() {
        Personnel.search(
                if      (state == "CA" || state == "MN");
                );
10     this.sort(last_name);
        this.format("%s,%s",state,last_name);
        Client.write(this);
        } main

15
```

GROUPBY: method

This method uses an embedded script within its arguments. Before the embedded script is interpreted by the method, the records that satisfy the request are selected and then sorted by the object(s) defined within parenthesis. The embedded script is then performed on the sorted records. Any variable values or objects changed in the embedded script will also be changed in the main script. The example below processes two states from the Personnel data object and then groups by state and returns two records containing the state and number of records for that state.

Example:

```
#import      Client
# import      Personnel
#local      recset      temp={ }
5          private     int       cnt=0 ;
main()  {
    Personnel.groupby( (state)
        if      (state == "CA" || state == "MN")
    {
        10      if (EQUAL)  {
            ++cnt;
        } else {
            temp.write(format ("state %s = %d",
                Personnel.state,cnt));
        15      cnt = 0 ;
        }
        }
    );
    Client.write (temp) ;
20      } //main
```

EXTRACT: method

This method uses an embedded script within its arguments. The
25 embedded script is interpreted by the method and executed for every record in the
method's object. Every record in the database that matches the search criteria will be
deleted from the data object, and a temporary object containing the records will be the
result of the method.

The example below will search the Personnel data object and return all records that state column is either "CA" or "MN". It also will delete these records from the data object. The records returned contain two columns (state and last_name).

Example:

```
5      #import      Client
          #import      Personnel

          main() {
10         Personnel.extract(
                  if      (state == "CA" || state == "MN");
                  );
                  this.format("%s,%s",state,last_name);
                  Client.write(this);
                  } //main
15
```

SAVEOBJECT: method

This method is used to create a new data object name or to create a data object on another server to allow parallel activity. The data object name identified by 20 the arguments is added to the list of system objects available to all members of the system. Once the object name is added to the system it requires a load or adon method to insert data. The method's object is an NT file identical to the result of the info method.

Example:

```
25      #import      Personnel
          main() {
                  Personnel.info().display().saveobject("objectx");
                  } //main.
```

MAIL: METHOD

This method is used to send e-mail. The method's object is sent as an attachment. There are three arguments of which only the first one is required. The first argument is a character string containing the person's e-mail address.

5 The example below searches the Personnel data object for employees with a particular birthdate. It then reads all records and sends e-mail to each employee with a birthday.

Example:

```

10          #import          Personnel
          #import          Birthdaycard
          private         record rec = new Personnel;
          private         char   buf[200];
          private string  bf="hope you have a wonderful day";
15          main ()  {
          Personnel.search(
          if      (birthdate == date$);
          );
          while(rec = this.read() )  {
          buf = format("Happy birthday %s %s",
          rec.first_name,buf);
          Birthdaycard.mail(rec.email,
          "birthday greetings",buf);
          } /while
          } /main

```

STATUS: method

This method is used to send single record character strings. There are different forms to the arguments.

5

```
Client.status("character string only");
```

```
Client.status(data_object);
```

10

The example below will send a character string, "We are performing your request".

Example:

```
#import Client
```

15

```
main() {  
    Client.status("We are performing your request");  
} //main
```

20

READ: method

This method suspends the script and waits until it receives an input.

The method loads the variable(s) in the arguments with the input. If there are multiple variables, then a comma in the input string will separate the data, and the variables will get loaded in order.

The example below will take a character string and parse it. The first part, until a comma is encountered, will be loaded into the bdate variable and the data following the comma will be interpreted as a number and loaded into the e_sal variable.

Example:

```
#import Personnel
# import Client
5    private char bdate[20];
    private int e_sal;

main() {
    Client.read( bdate, e_sal );
10   Personnel.search(
        if (birthday > bdate && salary > e_sal);

        )
    this.format("%s,%s,%s" ,last_name,first_name,ssn) ;
15   Client.write(this);
    } // main
```

20 SIZE: method

This method returns the number of records in the method's object.

The example below will get the number of records from Personnel data object and return the value.

Example:

```
5      #import      Client
      #import      Personnel
      private      int      records=);
10
      main()  {
      records = Personnel.size();
      Client.write(format ("The number of records in
      Personnel = %d",records));
      } //main
```

DISPLAY: method

15 This method displays records in a notepad window. Its purpose is to
help in the debug of a script.

The example below will open a notepad window with five records
from Personnel data object.

20 Example:

```
25
      #import      Personnel
      main()  {
      Personnel.display( 5 );
      } //main
```

COPY: method

This method loads the method's object into the data object in its arguments. If the column names of the objects don't match the columns are left 5 blank. Only the like columns are loaded. If columns are of the same type but different lengths, they will either be truncated or expanded to the size of the receiving column. If the method's object did not exist, it will be created. If the column types changed from the previous data object, then a remove method must be called before overloading the new object.

10

Example:

```
#import Personnel
#import Payroll
```

15

```
main() {
    Personnel.copy(Payroll);
} //main
```

20

ADTO: method

This method adds the method's object to the data object in its arguments. If the column names of the objects don't match the columns are left 25 blank. Only the like columns are added. If the columns are of the same type but different lengths, they will either be truncated or expanded to the size of the receiving column.

Example:

```
5      #import          Personnel
          #import          Payroll
main()  {
    Personnel.adto (Payroll) ;
} //main
```

10

ADON: method

This method adds the data object in its arguments to the method's object. If the column names of the objects don't match, the columns are left blank. Only the like columns are added if the columns are of the same type, but different 15 lengths, they will either be truncated or expanded to the size of the receiving column.

Example:

```
20      #import          Personnel
          #import          Payroll
main()  {
    Personnel.adon (Payroll) ;
} //main
```

25

REMOVE: method

This method removes the method's object from the system.

Example:

5

```
#import          Payroll

main()  {
    Personnel.remove();
} //main
```

10

MODIFY: method

This method uses an embedded script within its arguments. The
15 embedded script is interpreted by the method and executed for every record in the
method's object. This method updates every record in the method's object that match
the search criteria. Any variable values changed in the embedded script are also
changed in the main script.

The example below will add 10% to the salary column of the Personnel
20 data object where the state column is either "CA" or "MN".

Example:

```
#import Personnel  
  
5      main() {  
          Personnel.modify();  
  
          if (state == "CA" || state == "MN") {  
              salary += (salary * .10);  
  
10        }  
        );  
    } //main  
  
15
```

FIND: method

This method uses an embedded script within its arguments. The embedded script is interpreted by the method and executed on every record until it finds a record in the method's object that match the search criteria. The method then 20 inserts the record into a record variable for the main script to process.

The example below will find a record in the Personnel data object that match the search criteria and return it into the rec variable.

Example:

```
5      #import      Personnel
      private      record      rec= new Personnel 'main()
10      main()  {
      rec = Personnel.find (
      if      (ssn == "476504118")
      );
15      } //main
```

FINDLOCK: method

15 This method uses an embedded script within its arguments. The
embedded script is interpreted by the method and executed on every record until it
finds a record in the method's object that match the search criteria. The method then
will lock and insert the record into a record variable for the main script to process.

20 The example below will find a record in the Personnel data object that
matches the search criteria and then lock it. If the lock is available, it will load the rec
variable with the record. If the lock is not available, it will throw an exception.

Example:

```
#import          Personnel
private         record      rec= new Personnel
5
main()  {
    while(FALSE == (rec = Personnel.findlock (
        if      (ssn == "476504118")
10
        )))  {
            sleep(1000);
        } //while
    } //main
15
```

UPDATE: method

This method must be preceded by a findlock method in the same script as the method's object. The record object in the arguments is updated back into the 20 same record found in the preceding findlock.

The example below will find a record in the Personnel data object that matches the search criteria and return it into the rec variable. The phone column is changed and then the record is returned to the method's object.

Example:

```
#import          Personnel
private          record      rec  =  new  Personnel ;
5
main()  {
    rec = Personnel.findlock (
        if      (ssn == "476504118")
10
    );
    rec.phone = "425-881-5039";
    Personnel.update(rec) ;
} //main
15
```

COMPUTE: method

This method uses an embedded script within its arguments. The embedded script is interpreted by the method and executed for every record in the 20 method's object. Any variable values or objects changed in the embedded script are also changed in the main script.

The example below will only process two states from the Personnel data object and then group by state and return records containing the state and last name and first name.

25

Example:

```
#import          Client
# import          Personnel
#local      recset      temp = { }
5          local      recset      temp2 = { }
main ()  {
    Personnel.compute()

    if      (state == "CA" || state == "MN") {
        if (state == "CA") {
            temp.write(format("California %s %s",
10          Personnel.last_name, Personnel.first_name) );
            } else {
                temp2.write(format("Minnesota %s %s",
Personnel.last_name, Personnel.first_name));
            }
        }
15      }
    );
    temp.adon(temp2);
    Client.write(temp);
} //main
20
```

JOIN: method

This method creates a view of the data defined by the group of columns defined in the first set of parenthesis and then merges the columns from the 25 two records whenever the "if" part of the argument is true.

The example below will create a temp data object that contains three columns (last_name, first_name and city). The temp data object is then returned.

Example:

```
#import      Client
#import      Personnel
#import      Zipcodes
5
main() {
    Personnel.join(
        (last_name,first_name,Zipcodes.city)
        if (zip == Zipcodes.zip) ;
10
    );
    Client.write(this);
} //main
```

15

MPL: method (Multiple Points of Logic):

This command starts either a named script or a data object, which contains a script. If the method's object is a clustered object, then the script will be 20 started on each server with "this" set to the data object that resides on that server. If the method's object is not a cluster, then the script will be started on the same server or on a server specified by a second argument with "this" set to the method's object. When the scripts are completed the 'THIS' object from each server is returned and merged into a new 'THIS' result.

25

Example:

```
#import Personnel
#import Client
```

```
5      main() {
          Personnel.mpl("strip_data");
          Client.write(this);
      } //main
```

10 SAVESRIPT: method

This method is used to add the method's object (which is the script) and the script name identified by the arguments to the list of federation scripts available to all members of the federation. After the script is added to the federation 15 list of scripts, that script can be used as a script or as a method that applies to all data objects.

Example:

```
#local      nt      script=@scripts/example{ }
20      main() {
          script.display();
          script.savescript("scripta");
      } //main
```

PRINT: method

This method is used to send the method's object to a printer. If no arguments are supplied, the default printer is identified in the config file for this server. The example below searches the Personnel data object for employees with a 5 birthday this month. It will format the records and then print.

Example:

```
10 #import Personnel
    main() {
        Personnel.search(
            if (birthmonth == mon$ )
            );
            this.format("%s %s",first_name,last)_name);
            this.print();
    15 } // main
```

THIS: method

20 This method identifies the method's object as the current "THIS" object.

Example:

```
25 #import Payroll
    main() {
        Payroll.this();
        this.info().display();
    } //main
```

READ: method

This method reads a record from the method's object and inserts it into a record variable. A FALSE condition is returned when there are no more records.

5

Example:

10

```
#import Personnel
#local recset tmp-{ }
private record rec = new Personnel;

main() {
    while (rec = Personnel.read() ) {
        tmp.write(format("%s %s", rec.last_name,
                        rec.first_name) );
    } //while
    tmp.close() .display();
} //main
```

15

20

COLUMNS: command

This command creates a temporary data object that contains a list of all the columns of the data object.

The example will display using notepad a list of all columns in the data object.

25

Example:

```
#import Personnel
main() {
    Personnel.columns() .display();
} //main
```

METHODS: command

This command creates a temporary data object that contains a list of all the columns of the data object.

5 The example will display using notepad a list of all columns in the data object.

Example:

10 #import Personnel
 main() {
 Personnel.methods().display();
 } //main

15 The following code indicates how easily the system according to the preferred embodiment allows new functions to be invoked. In this case, a database, nt_data, is being searched for a particular State, and the results are being e-mailed (via Exchange) to the indicated address.

Example

20 # import nt_data

 main() {
 nt_data.search(if(state == "CA")).mail
 ("Bob.Hall@unisys.com")
 }

As another example, one may search a database containing the population of the United States for persons with particular attributes, sort the results of the search, and fax the results of the sort - in one easy statement. If the database

containing the population of the United States is distributed, the search and sort will be run automatically in parallel.

In installing the apparatus according to the preferred embodiment, one first defines a number of environmental variables. The value of FED_federation 5 number points to the “trader” for the federation. The trader contains a map network drive or network share. The drive contains the scripts, data object definition files errors and other messages for this federation. It also contains a file which identifies the servers in the federation. The latter is called the Federation file.

This and the following information is contained in a repository (acting 10 as a “trader”). The system administrator is guided, using a graphical user interface, through the specification of this configuration information. A repository application copies the information to the indicated local files (local to each server).

The value, FED_federation number.server number, identifies a base 15 directory for the server “server number” within the federation “federation number”. The apparatus uses the base directory for temporary files, scripts, data objects and NT files. The value of environmental variable, _PATH, provides the location of the apparatus executable.

The following illustrates a typical federation file, listing the servers 20 within the federation. (This file is contained within the “trader”, pointed to by FED_federation number.):

Federation Name = 1
Server = 10 “server name 10”
Server = 1 “server name 1”
25 Server = 2 “cluster 2”
Server = 3 “server name 3”
Server = 4 “server name 4”
Server = 5 “server name 5”

The following illustrates a typical server configuration file, contained within the base directory for the server, pointed to by “_FED federation.server.”

Typical configuration file

“Pre_start_processor” indicates the initial number of agent processes in the server. If additional processes are required, they will be generated dynamically. The “Inp.msg.time” parameter is a time out value.

The Debug parameter specifies levels of debugging information to be written to the trace file (the parameter has values from 1 to 10). The Sound parameter indicates whether audio error and other messages should be provided.

20 When the server is installed, the apparatus copies the federation file, all object definition files, and all scripts from the map network drive or network share (that is, from the “trader”) to the local subdirectories on the server. A “start service” is used to start or stop the servers in a federation.

25 Figures 10A-10C illustrates an alternate user interface in the form of a Java Studio based integrated development environment in which the apparatus (messenger, agent, assistant) is invoked as a Java Bean. A version of the apparatus invokable as a Java Bean is provided so that, in this case, the apparatus appears to be a Java Bean and may be invoked by, or invoke, other Java Beans.

The Java Studio™ contains a design panel (bottom right in Figure 10C) and a GUI panel (bottom left in Figure 10B). It supplies a pallet of Java Beans. When a Java Bean is dropped on the design panel, if it has a corresponding GUI element (e.g., the Java Bean might contain a GUI text field), the latter will appear in 5 the GUI panel (and may be moved, resized, etc.). When a bean is selected, there is the opportunity, via a dialogue box, to customize it. Thus, for example, a bean providing arithmetic functions might be customized with different operators, new expressions, etc., while GUI beans might be customized with respect to size, colors, background colors, fonts, etc., and other attributes.

10 Methods in the beans are exposed as connectors, so that the beans can be wired together visually. Components communicate by sending and receiving messages through the connectors. Messages may be a value (normal input, output); a trigger (e.g., a signal to start an action);...

Having created a design, one can generate an applet, standalone 15 application, java bean, or “packaged design” (a single representation of the whole design, which may be used in future designs). (An enhancement to the integrated development environment might show, for example, the web page into which the new applet – if generation of an applet were selected - may be dropped.)

The GUI is active, so that, as the design is built, the resulting solution 20 may also be tested. Key methods disclosed above (e.g., search, sort,...) are made available as Java Beans or ActiveX components, and are usable independently. The apparatus itself is available as a Java Bean or ActiveX component, driven by a simple script, and may be used to create applications (business objects) of any degree of complexity. The apparatus supports the construction of full mission critical 25 transaction or decision support solutions.

If the referenced data is distributed, the methods are invoked transparently, as parallel services – so that one may have transparent parallel data access and execution of methods across data distributed on the SMP nodes of a Unisys

PATENT DOCKET UNI6-BI30/04MV1089

Cellular Multiprocessor, cluster, or network. The transparently accessed data may be heterogeneous: contained, for example, in NT files, SQL Server, Oracle,...

In the diagram of Figures 11A-11D, a search bean is shown driven by the indicated script. A "Personnel" database is imported, and is searched – with the 5 command Personnel.search – for persons with a particular birthdate. If the "Personnel" database is distributed, the search methods will be run, automatically in parallel, across the distributed data. The search bean is easily customized by modifying the script. The example shows the supplied bean(s) being invoked by, and invoking, beans provided by Java Studio or by other bean "factories."

10 The considerable increase in performance of a system constructed according to the preferred embodiment may be illustrated in connection with an actual example concerning mining of the population of the United States. In this example, five million records, spread across 15 nodes, were searched, in serial, in 3 minutes, 27 seconds. Using a parallel search according to the preferred embodiment, the time was 15 17 seconds. If indices are built for the attributes of interest and searched in parallel, the search time is 4 seconds. If the indices are then cached, and the cached indices are searched in parallel, the search time is 1.2 seconds (some 172 times faster than the original serial search).

20 The times will often be linear, so that ten million records on 30 nodes, or 2,500,000 records on 8 nodes, will be searched in the same 1.2 seconds. The ability to search the population of the United States in less than 16 seconds portends a tremendous opportunity for completely new applications and solutions.

Another example of increased performance is a search of 15 million records across five nodes is illustrated in the following table:

		<u>Serial</u>	<u>Parallel</u>	<u>Improvement</u>
		(preferred embodiment)		<u>Times</u>
NT file system	-search	6 min, 32 sec.	1 min. 22 sec.	4.44
	-sort	26 min, 26 sec.	6 min., 8 sec.	4.31
	-indexed		3 sec.	130
	search			
SQL Server	-search	6 min., 11 sec.	1 min., 20 sec.	4.64
	-sort	2 hr., 38 sec.	25 min., 20	4.76
			sec.	
Oracle	-search	14 min., 5 sec.	3 min., 8 sec.	4.49
	-sort	3 hrs., 17 min.,	36 min., 18	5.45
		43 sec.	sec.	

5 In the above example, servers were installed on five Aquanta systems, connected to a PC client with 10T Ethernet. The installation of the servers transformed the environment into a cluster or multiprocessor federation. In this way, the servers may be used, for example, to support a virtual data warehouse (the data in which may reside in heterogeneous databases where a database need not scale beyond 10 the largest server).

10 As noted above, the preferred embodiment employs a function shipping model. An alternative model is a data shipping model in which, rather than the functions being performed as close to the data as possible, and only the results (e.g., the number of "hits" in a search) being returned to the requester, the data to be 15 searched is passed to the requester and the function is performed there. This involves

the transfer of much more data about the network, and, even with very high network or backplane bandwidths, typically increases the latency of the operations. The implemented model also ensures that the latest version of the data in question never resides in a cache on a different node (thus eliminating cache coherency problems).

5 Certain design considerations may also be noted at this point. Critical
in terms of scaling and performance are concurrency (the degree of parallelism),
contention (the degree of serialization) and coherency (the cost of consistency). An
aim is to minimize contention and the cost of consistency. While one cannot totally
eliminate contention, it is found that by performing updates as close to the data as
10 possible (treating that server as largely independent of the others), and by
randomizing the distribution of data across the servers (e.g., using hashing
techniques), the contention is usually low (shown in almost linear scaling).

The cost of coherency (not having a later update in someone else's cache) is a quadratic term, placing a maximum on the possible scaling as servers are added. By performing updates on the server containing the data, one ensures that there will never be later updates in the cache on other servers, ensuring cache coherency, and eliminating this cause of a maximum in the possible scaling.

Additionally, the preferred embodiments discussed above have employed the creation of data descriptor files from the metadata at run-time. In alternate embodiments, all files can be held in the repository, which is then accessed from each server at run-time or a distributed repository system may be used where lightweight versions of the repository reside on the local server nodes.

The methods and apparatus of the present invention, or certain aspects or portions thereof, may take the form of program code (i.e., instructions) embodied in tangible media, such as floppy diskettes, CD-ROMS, hard drives, or any other machine-readable storage medium, wherein, when the program code is loaded into and executed by a machine, such as a computer, the machine becomes an apparatus for practicing the invention. The methods and apparatus of the present invention may also be embodied in the form of program code that is transmitted over some

transmission medium, such as over electrical wiring or cabling, through fiber optics, or via any other form of transmission, wherein, when the program code is received and loaded into and executed by a machine, such as a computer, the machine becomes an apparatus for practicing the invention. When implemented on a general-purpose 5 processor, the program code combines with the processor to provide a unique apparatus that operates analogously to specific logic circuits.

Those skilled in the art will appreciate that various adaptations and modifications of the just-described preferred embodiments can be configured without departing from the scope and spirit of the invention. Therefore, it is to be understood 10 that within the scope of the appended claims, the invention may be practiced other than as specifically described herein.